BUILT BY A TEAM THAT HAS SECURED MAJOR FINANCIAL INSTITUTIONS WORLDWIDE FOR 25 YEARS.

ARGENT FOR BANKS

AI AGENTS NOW ACT LIKE STAFF. PUT A CHECKPOINT IN FRONT OF EVERY ONE.

Argent Strategic Intelligence puts a checkpoint between your AI agents and the outside AI services they use. It checks every request before it leaves, and it runs entirely inside your own bank. Nearly half of banks are testing AI agents,[1] and FINRA's 2026 report added agents acting beyond their authority as a new risk.[2]

One pane of glass for your entire agent ecosystem: every agent, model, tool, owner and dollar in one view. **No rip and replace.** Argent works with what you already run: your identity provider (Okta, Microsoft Entra, LDAP), your own secrets vault (HashiCorp Vault, AWS Secrets Manager, CyberArk Conjur) and the AI tools your staff already use.



SEE.
Every agent, its owner, where it sends data.



PROVE.
A signed record of every decision.



STOP.
Unapproved models, tools and leaks are refused.

THE FIVE THINGS BANKS ASK FOR FIRST

1	**Budget controls** A trading desk's agent runs up a large bill overnight.	Set spend caps by team, person and agent. An agent that hits its cap is stopped; its owner asks for more, and nothing is granted until a manager approves, in Argent or through your ServiceNow or Jira tickets. Project cost tracking shows what each project or business line spends on AI.
2	**Data protection** A loan officer pastes an applicant's whole file into an AI.	Argent reads each request first and blanks out Social Security, card and bank account numbers before anything leaves the bank (routing and SWIFT codes too, where you switch them on). The record notes what kind of data, never the data itself.
3	**Auditing for auditors** An examiner asks which agent sent what, and when.	Every AI decision goes into a signed, tamper-evident record that cannot be quietly changed. An examiner or internal audit can check a package of it on their own computer and see whether anything was altered.
4	**Model control** Staff start using an AI model nobody approved.	Only models you approve are allowed; everything else is refused. Group them into model classes, such as approved for customer data and general use. A model can be swapped for an approved one without staff changing anything.
5	**Air-gapped** Customer data must never leave the building.	Argent runs entirely inside your own bank, even on a fully air-gapped network. The install bundle is checked at every step as it crosses into your closed environment.

WHAT REGULATORS ALREADY EXPECT

Federal banking agencies. A bank must tell its primary federal regulator about a serious computer-security incident within **36** hours of deciding it is one.[3]

New York DFS, 23 NYCRR 500.17. Notice "as promptly as possible but in no event later than **72** hours" after determining an incident has occurred.[4]

HOW A BANK STARTS

1. Watch mode. A real-time audit and view of what your AI is doing today. Nothing is blocked.

2. Enforce mode. Switch on blocking one area at a time: customer data, budgets, models, then tools.

3. Inside your walls. Runs on your own hardware or in your own cloud VPC. Nothing leaves the perimeter without your explicit approval.

BOOK A DEMO. WE WILL SHOW YOU ARGENT WORKING ON YOUR BANK'S OWN USE CASES.

NYEDIS.AI/DEMO

Sources. 1 KPMG LLP, AI Quarterly Pulse Survey: Banking, Q1 2026, April 2026. 2 FINRA, 2026 Annual Regulatory Oversight Report, GenAI: Agents, Dec 2025. 3 OCC Bulletin 2021-55, Computer-Security Incident Notification Rule (86 FR 66424), Nov 2021. 4 23 NYCRR 500.17(a)(1), NYS DFS, read 2026-11-18.